

(58) Field of Search
UK CL (Edition S) H4P PEUX
INT CL⁷ G06F 11/34 , H04L 12/24 12/26 29/06
Online: EPODOC, JAPIQ, WPI

Fig.2





Fig.1

**Fig.2**

3/5

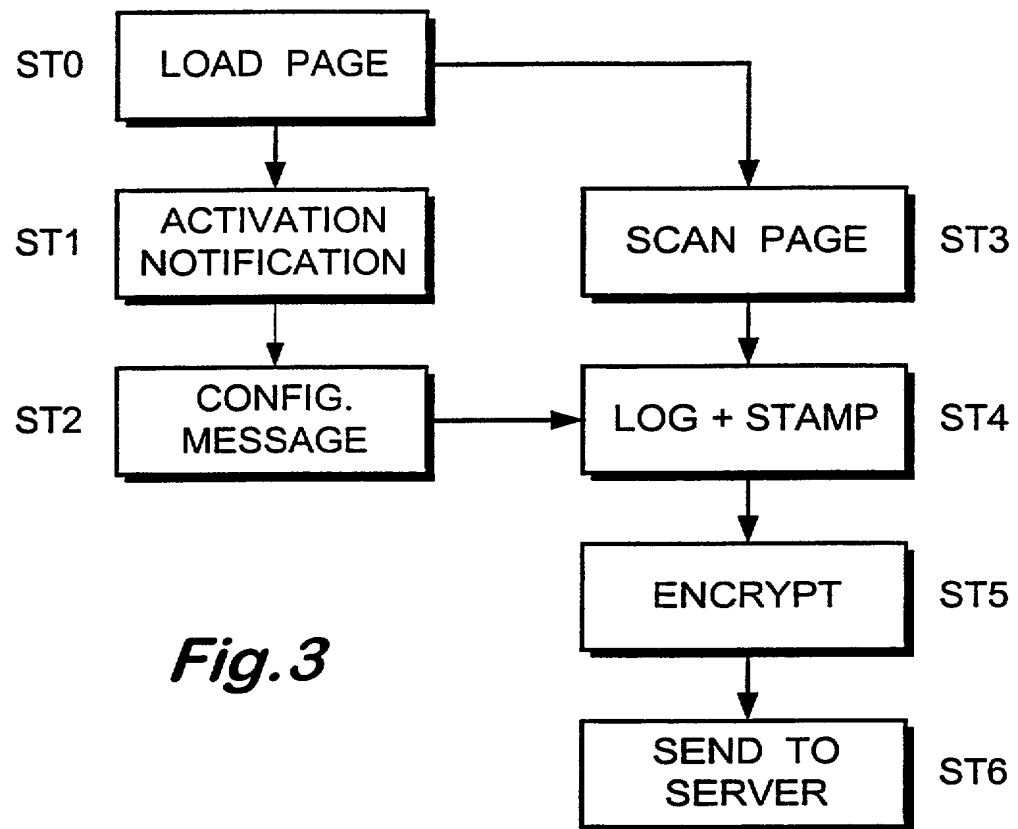


Fig. 3




Fig. 4

Fig. 5

IndexNo	LoggedTimeStamp	EventTimeStamp	OriginatorID	RecordType	Context
29	06/02/00 14:38	06/02/00 14:38	169.254.1.2	Initialise	0
30	06/02/00 14:38	06/02/00 14:38	169.254.1.2	Navigate	http://www.microsoft.com/
34	06/02/00 14:46	06/02/00 14:46	169.254.1.2	Initialise	0
35	06/02/00 14:46	06/02/00 14:46	169.254.1.2	Navigate	http://www.microsoft.com/
36	06/02/00 14:47	06/02/00 14:47	169.254.1.2	DocLoad	[http://www.microsoft.com/]<HTML><HEAD><<
37	06/02/00 14:47	06/02/00 14:47	169.254.1.2	DocLoad	US/Worldwide></
38	06/02/00 14:49	06/02/00 14:49	169.254.1.2	Initialise	0
39	06/02/00 14:49	06/02/00 14:49	169.254.1.2	Navigate	http://www.speed-trap.com/
40	06/02/00 14:50	06/02/00 14:50	169.254.1.2	DocLoad	[http://www.speed-trap.com/]<HTML><HEAD>
41	06/02/00 14:53	06/02/00 14:53	169.254.1.2	Navigate	http://www.microsoft.com/

5/5

Fig.6

Monitoring User Interactions with Services provided
over a Network

This invention relates to monitoring user
5 interactions with services provided over a network.
In particular it relates to monitoring interactions
with services provided over an intranet and/or the
Internet and particularly the World Wide Web.

There is currently great interest and activity in
10 providing services via networks and particularly the
World Wide Web. However, some difficulties arise for
providers and users in providing and using such
services. In particular, waiting for web pages to
download from the World Wide Web can be time consuming
15 and this factor together with a poorly designed
website can severely limit the efficiency of web based
services. Whilst these difficulties with web based
services are widely appreciated, there have, to date,
been no satisfactory systems available for gaining
20 data concerning the effectiveness and usability of any
given web based service. Existing systems tend to
provide generalised information by monitoring networks
and servers.

It has been appreciated that obtaining detailed
25 information concerning the use made of websites by
individual users can be invaluable in enhancing the

services which are provided over the Internet and particularly the World Wide Web.

It is an object of this invention to provide a system for monitoring the operation of services
5 provided over a network. At least some embodiments are directed to obtaining information concerning how the performance of the service is perceived by users.

According to a first aspect of the present invention there is provided a method for monitoring
10 user interactions with a service provided over a network to a plurality of end user communications units comprising the steps of:
monitoring and recording activities instigated by at least one user when using the service by interacting
15 with one or more electronic page displayed in a browser on a respective communications unit, the recorded activities including a log of pointing device position on the or each page; and
playing back activities recorded in the above step,
20 wherein the playing back step includes redisplaying the electronic pages originally displayed and redisplaying user interaction therewith.

According to a second aspect of the present invention there is provided apparatus for monitoring
25 user interactions with a service provided over a

network to a plurality of end user communications units comprising means for monitoring and recording activities instigated by at least one user when using the service by interacting with one or more electronic page displayed in a browser on a respective
5 communications unit, the recorded activities including a log of pointing device position on the or each page; and
means for playing back the recorded activities by
10 redisplaying the electronic pages originally displayed and redisplaying user interaction therewith.

Such arrangements make use of the aggregate computing power available at the many user communications units which are typically connected to
15 a network rather than attempting to support such data gathering with a single or a few central servers. It is also possible to collect data which is specific to the experience of each of a number of users.

It will be appreciated that playing back of
20 recorded activities refers to the apparatus itself, under the control of suitable software, running through any and all steps instigated by a user when the recording was carried out.

In developments of the invention, activity
25 associated with the service may be monitored during

the play back of the recorded activities. This monitoring may be carried out in substantially the same way as defined in any of the other aspects of the invention defined above. In particular, the activity
5 may be monitored at each communications unit and sent to the central unit.

In this application "activity"/ "activity instigated by the user" includes actions directly carried out by the user such as "clicking" of a mouse
10 on an icon and actions carried out by the system without the user being directly involved, for example, data being sent and/or received. In practice activity will typically comprise a stream of occurrences consisting of actions taken by the user and events
15 taking place within communications unit.

The methods may comprise the further step of analysing the data obtained during the monitoring step. The apparatus may be arranged for analysing monitored data. The analysis may comprise filtering
20 the data so that it can be presented in a more convenient form. One or more filter rules may be specified for use in the filter process. The or each filter rule may define a standard event which comprises a predetermined set of occurrences. One or
25 more filter rule may specify at least one parameter,

values for which are to be extracted from the predetermined set of occurrences. The filtered data may be presented in terms of standard events and associated values. Preferably the analysis is carried out at the central unit.

The term end user communications unit is used broadly in this application and should be taken to include conventional computers including PCs, interactive television based systems, landline and mobile communications (telephone) based systems and any other such system which is able to receive information provided across a network. The system may be used with units operating in accordance with WAP (Wireless Application Protocol).

The central unit to which monitoring information is sent is generally distinct from any server supporting services which are to be monitored. The central unit may belong to an organisation which is interested in gathering its own performance information or to an organisation administering the performance monitoring system on behalf of customers. Information collated at the central unit may be accessible via a web site. Customers may be allowed access to information at the central unit which relates to the services which they supply. The central

unit may comprise a plurality of devices, for example servers, operating in parallel. In some embodiments more than one central unit may be provided.

The network over which services are provided typically comprises an intranet and/or the Internet. The services provided are typically accessible via the World Wide Web. Thus the services are typically accessible via a web site. The browser may be a generally available browser such as MS Internet Explorer or Netscape Navigator^(RPM). Alternatively a dedicated browser may be provided. The browser may be an embedded application in, for example, a WAP communications unit.

Any one or any combination of the monitoring, outputting, recording and play back functionality at the communications unit may be facilitated by an applet. Preferably the applet runs in a language which is cross-platform. This facilitates the system functioning on a large range of different end user communications units. The applet may be confined in an environment such that the applet's access to the end user's communications unit is restricted. This can significantly enhance security.

In this application the term applet is used broadly to mean an application which runs within a

browser. The applet may be in any appropriate language, for example Java^(RTM) and/or JavaScript^(RTM). However, use of the term applet should not be taken to specifically imply that the applet is written in/runs in Java^(RTM). The applet may be persistent or non-persistent. Thus, the applet may be resident in the browser at all times so that activity relating to all pages is monitored or the applet may be associated only with certain web pages so that only activity in relation to such pages is monitored. In the persistent case the applet is installed, at the communications unit, by a user or system administrator. In the non-persistent case the applet is loaded and deleted, at the communications unit, as a page is visited and left respectively.

The communications unit may be arranged so that as a page is visited and loaded into the unit, the interactive content of the page is identified. For example, a page may include a button which can be "pressed" by a user and field allowing the input of text. In the identifying step, the unit may identify that the button and text field are present. If the button is pressed or text is entered this action may be monitored and recorded or reported.

The identifying step may include determining

information concerning the relevance and/or context of the interactive content. For example, the relevance, context, effect etc of a user entering text or pressing the button may be identified. This
5 information can be associated with any recorded or reported action.

In at least some embodiments any one or any combination of the following may be monitored: start of page load, completion of page load, completion of
10 image or frame load, user input of data starts, user input of data finishes, user data, hot spot and hyperlink roll-overs, hyperlink selection, mouse position, field selections, browser toolbar usage, next Uniform Resource Locator (URL) selected, user
15 editing field information, user identification (log on name, device, Internet Protocol (IP) address etc.).

The data output and/or sent to a central unit as a result of the monitoring process may be encrypted. Since it is operation characteristics which are of
20 interest, certain sensitive data, such as credit card details, may be irrelevant. Thus any such sensitive data may be replaced by arbitrary symbols in the information output and/or sent to a central unit.

Timing means may be provided for giving time
25 stamp references to activities. Preferably a universal

time reference source is provided at the central unit.

In at least some embodiments any one or any combination of the following real time displays or reports may be generated for a site being monitored:

5 number of visitors per unit time, number of new customers, number of customers who ordered, number of customers by location, number of visitors who tried and failed to order, total web site orders, value of total web site orders, usability of web site, apparent
10 speed per customer, typical routes through site, page hot spots or trouble spots, time to complete transaction pair, events by time of day, average/minimum/maximum time to complete transaction, page load time by time of day, time from page load to
15 page exit, exit route from each page, name of user by time of day, IP address of user by time of day.

The or each communications unit may, amongst other things, function as a data capture module. The central unit may, amongst other things, function as a
20 database server.

According to a third aspect of the present invention there is provided a machine readable medium having thereon code portions which when loaded and run on computer means cause the computer means to execute
25 any of the methods defined above and/or provides any

of the apparatus defined above.

Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings in which:-

Figure 1 schematically shows the architecture of a system for monitoring the performance of services provided over a network;

5 Figure 2 schematically shows a system for monitoring the performance of a website based service provided over the World Wide Web wherein the monitoring activity takes place within a browser;

10 Figure 3 schematically shows the steps carried out in monitoring activity associated with a web page in the system shown in Figure 2;

Figure 4 schematically shows a process used to test a server system;

15 Figure 5 schematically shows internal architecture of a database server; and

Figure 6 shows an example of data stored in a data log of the database server shown in Figure 5.

20 Figure 1 schematically shows the architecture of a system for monitoring the performance of services provided over a network.

25 A plurality of end user communication units 1 which can include conventional computers, interactive television based systems and wireless communication systems etc., are connected to an intranet or the Internet network 2 and thus are able to access

services supported by a large number of different servers 3 (only one of which is shown in Figure 1). A central unit 4 to which data concerning the performance of services provided over the network is to be sent, as will be explained in more detail below, is also connected to the intranet or Internet 2. A customer's computer system 5 is connected via the intranet or Internet 2 to the central unit 4 to allow the customer to extract information concerning the performance of services in which he is interested. Although not shown in detail it will be appreciated by those skilled in the art that within the intranet or Internet 2 there will be a number of intermediate devices such as routers, servers and other devices. These are all potentially connectable to the central unit 4 and as such data concerning performance of these intermediate devices may also be monitored and sent to the central unit 4. A simplified description of the functioning of the system is given below.

At each of the end user units 1 a user will access a service supported by a web server 3 via the intranet or Internet 2. As this service is used by the user then, under the control of appropriate software, each end user unit 1 monitors activities instigated by the user which relate to the service

being used. The data obtained during this monitoring step is then sent from each of the end user communication units 1 via the intranet or Internet 2 to the central unit 4. The central unit 4 functions as a database server and performs a number of functions as will be described in more detail below. At this point it is enough to note that all of the data obtained at each of the end user communication units 1 is stored at the central unit 4 and is put into a useful form such that a customer may obtain data concerning the services in which he is interested either in the form of hard copy reports or via his computer system 5 and the intranet or Internet 2. In particular it is envisaged that the customer will operate a web site based service and will desire information concerning the performance of that website.

Each of the end user communication units 1 includes a processor, a memory, and other elements which are arranged under the control of appropriate software to capture data concerning various activities which take place in respect of the service being provided over the network 2. Different pieces of software are provided to give this data capture function in different circumstances. Different pieces

of such software will be described in more detail below.

As mentioned above data capture modules can be arranged to collect information from a variety of sources. In particular data can be collected by monitoring activities within a user's browser, from the network delivering data to the user, and from servers hosting the web pages or applications which a user is using and/or running. However, it will be appreciated by those skilled in the art that the collection of data need not necessarily be limited to the above-mentioned sources.

The central unit 4 includes a processor, a memory and other elements which are arranged under the control of appropriate software to act as a database server which is used for receiving, storing, analysing and organising data captured by any of the data capture modules as well as generating reports concerning such data.

Appropriate software is also provided to customers for use with their computer systems in order to further process and interpret the data stored at the central unit 4. In particular it is possible for customers to generate reports, analyse statistics and generate graphs on the basis of the collected

data.

It should be noted that the methods and apparatus for monitoring and/or testing the performance of services provided over a network which this application concerns may be used in relation to any appropriate service provided over a network. The methods and apparatus are not limited to use with services provided via websites accessible over the World Wide Web. However, for the sake of clarity the remaining detailed description will be phrased in terms exclusively relating to services that are accessible via a website and the World Wide Web.

Figure 2 schematically shows a system for monitoring the performance of a website based service provided over the World Wide Web wherein the monitoring activity takes place within a browser such as Netscape Navigator or MS Internet Explorer loaded on a user's computer system which acts as a communications unit 1.

In the case shown in Figure 2 the data capture module consists of a user's computer system 1 arranged under the control of a browser and an applet which is resident in each web page that is to be monitored. Applets are small applications written in, for example, the Java^(RPM) and JavaScript^(RPM) programming languages

which may be downloaded from a website into a user's browser in order to carry out a particular function. The applet used in the system shown in Figure 2 is currently named Prophet by the applicants and this
5 name shall be used in the description to aid understanding.

It will be appreciated that a system such as that shown in Figure 2 can be used to monitor performance of a large number of websites provided at different
10 servers which are connected to the Internet. However, again for the sake of simplicity in this case, a single website and the performance of the service supported by that website shall be considered.

Thus, in this case, a website is supported by a
15 web server 3 which is accessible via the Internet 2. The website has been set up such that the Prophet applet is contained by each web page 101 which is to be monitored. This means that as a user visits a web page 101 the user's browser automatically downloads
20 and executes the applet. Figure 2 schematically shows web pages 101 displayed in the user's browser and the prophet applet 102 running within the page 101.

The user's communication unit 1 under the control of the prophet applet monitors the user's interaction
25 with the web page 101 and returns the information

obtained to the central unit 4 which in Figure 2 is shown to comprise a database server 41. A customer can obtain information from the database server 41 concerning the performance of the website provided at the web server 3.

Each frame of each web page to be monitored must include a copy of the Prophet applet.

The Prophet applet may be configured to perform a range of different monitoring tasks and to select different types of data. In practice a single version of the prophet applet is made available and included into the appropriate web pages by the author or maintainer of the pages. The configuration is achieved by a configuration message which is sent to the applet from the central unit server 41. Figure 3 is a flow chart schematically showing the steps carried out in monitoring activity associated with a web page in the system shown in Figure 2.

In step 0 a web page and resident Prophet applet is loaded into a user's browser. In step 1 as the applet is activated within the browser an activation notification signal is output under the control of the applet to the central unit server 41.

Each customer (each entity who includes the Prophet applet on their web pages) is supplied with a

registration key. This registration key is used to verify web page 101 and central server 41 match. The appropriate key is included in the applet in each page. A copy of the registration key is included in
5 the activation notification. The central server 41 checks to see that a valid key has been received. If there is some mis-match this could indicate an incorrectly set up page or a hacking attempt. If any such mis-match is determined then any further data
10 received from that instance of the applet is ignored by the central server 41.

The activation notification includes an indication of the instant at which the notification is generated in terms of time as kept by the
15 communication unit 1 concerned. The Prophet applet treats this instant as time zero and all later occurrences are logged relative to time zero. The time of this instant, in communication unit time, is received by the central server 41 and compared to a
20 server universal time. The result of the comparison is used to apply an appropriate offset to subsequent timings received from the communications unit 1 such that all occurrences dealt with at the central server 41 are logged in server universal time.

25 Correct receipt of the activation notification by

the central server 41 triggers issuance of a configuration message, at step 2, which is sent to the respective communications unit 1. The content of the configuration message is selected by the central unit
5 server 41 on the basis of preselected information received by the server 41. In particular, the configuration will depend on the page being loaded into the browser and may also depend on other factors or information such as the IP (Internet Protocol)
10 address or name of the user, the time of day and/or other user or system information.

The configuration message includes a session ID which can be used to identify the origin of each occurrence which is monitored and sent to the central
15 unit server 41. This session ID uniquely identifies the particular browser being used and the particular page being used so that occurrences relating to a single page can be tied together easily. There is a further advantage in that, if a page is downloaded
20 twice, two different session ID's will be generated. This fact is detectable in the reported data and can indicate some kind of problem such as abnormal reloading behaviour or partial loads taking place because of an annoyingly slow network etc.

25 At the same time that the applet is signalling

its existence in step 1, the contents of the page being loaded into the browser is scanned by the prophet applet in step 3. This scanning step is facilitated by a suitable JavaScript^(KTA) process. In particular the prophet applet operates by interpreting the HTML of the page concerned as it is loaded and from this it develops a page map which includes a list of all of the page components. This is used to allow the applet to quickly analyse any action taken by the user and prepare a record which correctly contextualises the action.

At step 4 a log of occurrences which occur in relation to the web page concerned is generated, each occurrence being given a time stamp relative to zero time. The occurrences logged in step 4 correspond to actions directly taken by the user as well as events occurring in the browser without the user's direct involvement.

A large range of different occurrences may be logged. As mentioned above the actual occurrences logged will depend on the contents of the configuration message received by the prophet applet. Below is a non-exhaustive list of the occurrences which may be monitored and logged:

Start of page load

Completion of page load

Completion of image or frame load

User input data finished

User input data started

5 User data

Hot Spot, and Hyper Link rollovers

Hyper Link selection

Mouse position

Field selections

10 Browser toolbar usage

Next URL selected

User editing of field information

User identification (log on name, device, IP address, etc)

15 In step 5 any sensitive information such as credit card details may be replaced by non specific characters such as asterisks and all data is encrypted.

20 At step 6, the information logged at step 4 and encrypted at step 5 is sent via the Internet 2 to the central server 41 for further processing and reporting as will be described in more detail below. In practice data will be allowed to build up to a predetermined level at the communications unit 1 before being sent
25 to the central server 41. Steps 4 to 6 in Figure 3

will run and re-run continually whilst a page is active but it is not necessary for all three steps to be carried out for a given occurrence before processing of the next occurrence begins. An
5 additional sending of data may be made as a page is left to ensure all data has been output.

As has been discussed above, such a system allows data to be collected which gives an insight to user behaviour and user perceived performance in relation
10 to particular web pages. This information can be used to help to improve the performance or effectiveness of services provided via web pages.

An alternative data capture module is implemented by use of a browser plug-in currently called Vigilante
15 by the applicants which functions in a similar way to the prophet applet but is permanently resident in a user's browser. Thus, in a Vigilante based system, rather than the applet being downloaded onto a user's communication unit each time a web page is loaded,
20 Vigilante is always present in the user's browser and thus can be used to monitor a user's interaction with any web pages visited. This can be achieved either by the Vigilante plug-in directly monitoring a user's activity or by it being arranged to load the prophet
25 applet into each downloaded page. The Vigilante plug-

in is typically included in the user's browser by the user or the user's system administrator. The Vigilante plug-in may be used for monitoring activity in a similar way to as the Prophet applet discussed above but also may be used for security purposes. For example, an Internet Service Provider may offer an incentive to its users to use Vigilante whilst logged on through their service or an employer may use it to monitor activities being carried out by his or her employees. Although the terms applet and plug-in are used herein to aid understanding, it should be noted that when used in its broad sense, the term applet encompasses plug-ins.

As a further alternative a dedicated browser may be provided in which the Vigilante or prophet functionality is incorporated. The dedicated browser may be used in substantially the same way as existing browsers such as Netscape^(RPM) Navigator and MS Internet Explorer, whilst at the same time providing the monitoring functions discussed above. The browser might be used by a user or a customer to monitor and/or control a user's activity. This may be useful for organisations to retain a certain level of control over staff's access to the intranet or Internet or at least monitor their activities. In such systems it

may be preferable to provide a configuration signal from a local server or stored cookie rather than from the central unit server 41.

5 Data concerning network and server operation may also be monitored by the provision of suitable software. The exact method for monitoring such data will depend on the devices being monitored. However, the use of a configuration message and the use of server universal time is preferable. Further, the data
10 extracted is preferably returned to a central unit for further processing.

In a further developments, a server supported service testing and in particular a load testing tool is provided. The testing system uses the same general
15 architecture as is shown in Figure 1. Figure 4 schematically shows the process used to test the server based system.

Initially one or more users is allowed to use a service supported by the server 3 which is to be
20 tested, and whilst doing so their activity is monitored in step 100 using the prophet or Vigilante systems described above. This monitored activity is then recorded as a script in step 101. In step 102 one or more scripts recorded in step 101 are replayed
25 by end user communication units 1. Typically a large

number of scripts will be run on separate end user communications units 1 to simulate the effect of the service experiencing heavy usage. In step 103, as the scripts are replayed the performance of the service is
5 monitored from the point of view of each user, ie at each end user communications unit 1.

The second monitoring step 103 is again substantially the same as that described above in relation to the Prophet applet and Figures 2 and 3.
10 At step 104 the data gained in the monitoring step is sent to the central unit 4.

Suitable software for controlling the running of scripts and monitoring of behaviour is provided at each communications unit 1. Although this may be
15 implemented in a number of different ways it will be most efficient to use applets or plug-ins substantially the same as the Prophet applet and Vigilante plug-in described in detail above. The software used to implement steps 102 to 104 shown in
20 Figure 4 above is currently called Slayer by the applicants and this term will be used below to ease understanding. When Slayer is run on a end user communication unit 1 it announces its presence to the central unit server 41 and awaits further
25 instructions. Further instructions are issued via the

central server 41 typically under the control of a systems manager in charge of the server based service to be tested. The necessary user communications units 1 are activated, for example by a systems manager or
5 an automated system, and caused to load a particular web page. Once this has been done interaction with web pages is governed by the recorded scripts which may be downloaded from the central server 41 to the user communication units 1 concerned.

10 It is to be noted that the scripts to be replayed by Slayer can, and preferably should be recordings of activities carried out by actual users. Thus recorded scripts are likely to include mistakes, incorrect entries, invalid transactions and perhaps arbitrary
15 delays or hesitations. This can give a much more accurate representation of real world traffic.

However, although the scripts are typically recordings of real users and can be used to accurately re-create the actual steps taken, the scripts may be
20 replayed at high speed and/or multiplied to give the effect of many users simultaneously attempting to use the same system. All factors concerning the rate at which data is entered and the time between each transaction etc. may be controlled by the system
25 manager to give the desired test circumstances.

A further advantage of this system over one which relies on artificially generated test programs is that when web page designs are changed, scripts relating to the amended pages can be automatically generated as users use those new pages. If a Prophet or Vigilante based system is being used in relation to the pages concerned there will be a ready made source for recorded scripts. The need for manually re-generating scripts for each version of the site can therefore be avoided.

The following are examples of browser events which may occur as the user interacts with a page. In each case these events may be associated with scripts. This means that such events may be monitored using the Prophet or Vigilante system mentioned above and also may be re-run using the Slayer system.

Load:- occurs when all the content of a particular document has been loaded within a page.

Unload:- occurs when a document is removed from a window or frame.

Abort:- occurs when page loading is stopped before an image has been allowed to completely load.

Error:- occurs when an image does not load properly or when an error occurs during script execution.

Select:- occurs when a user selects some text in the

text field.

Focus:- occurs when an element receives focus via a pointing device or tab navigation.

Change:- occurs when part of a page such as an input
5 region or a text area loses focus and its content has been changed since gaining focus.

Submit:- occurs when a form is submitted.

Reset:- occurs when a form is reset.

Blur:- occurs when an element loses focus.

10 Resize:- occurs when a document view is resized.

Scroll:- occurs when a document view is scrolled.

The arrangement of the present systems give advantages in that aggregate computing power of a large number of end user communication units may be
15 used in the gathering of data and or the running of scripts to provide a load testing function. This limits the amount of processing power required at the central server 41 and facilitates the gathering of information which is specific to individual end users.

20 To maximise the advantage of this distributed processing effect the format of the data sent to the central server 41 should be such that as little processing as possible is required on the data received at the server.

25 Although the system is such that data can only be

gathered in relation to pages in which the applet is active, it is possible to gain information about what happens to cause the page to be exited. Typically this may be movement to another page or website. The
5 applet may monitor and log the identity of the next site to be visited and also the event which caused such a move, for example a mouse click on a hypertext link or the direct entry of a URL.

In some circumstances it may be necessary for the
10 applet to be signed using a signature technique which is applicable across all compatible browsers. The authority granted by the certificate should be sufficient to accomplish tasks required for monitoring and logging data. Thus the certificate is preferably
15 sufficient to allow access to all entities within the HTML frame, to obtain access to sub-windows opened by the browser or any contained scripts, and optionally to access the browser buttons and menu activity. If a signed applet is received by an end user communication
20 unit 1 this version of the applet may be retained by the end user communication unit and used for subsequent pages.

Figure 5 schematically shows the internal architecture of the database server 41 supported by
25 the central unit 4. The database server comprises a

data packet receiver/generator 6 which is arranged to output and receive data to and from any data capture module. The data packet receiver/generator 6 facilitates communications between the data capture
5 modules and the database server 41. Amongst other functions, the data packet receiver/generator 6 receives activation notices from data capture modules and returns configuration messages.

The database server 41 further comprises a system
10 maintenance module 7 which provides tools and mechanisms for defining configuration rules for the different versions of data capture modules which are in existence. In relation to the Prophet applet this allows the generation of configuration settings for
15 use with the different instances of the Prophet applet which have been included in web pages by customers.

The tools provided by the systems maintenance module 7 also allow systems managers to populate and manage an originator's database 8 which holds
20 configuration data for all known data capture modules, i.e. all known versions of the Prophet applet amongst other things. The originator's database 8 constructs and issues configuration messages based on the activation notifications which it receives from remote
25 data capture modules via the data packet receiver 6.

The system maintenance module 7 also provides configuration and maintenance tools for the other functions of the database server 41 described below.

The database server 41 includes a data log 9
5 which holds all the data collected from the data
capture modules. Each entry in the data log includes
details of the occurrence and a time stamp indicating
when the occurrence was lodged in the data capture
module. The data log 9 also adds a time stamp
10 indicating the time at which it logged the occurrence.
Figure 6 shows an example of the data stored in the
data log 9.

As well as data from user browser based data
capture modules, the data log 9 also logs utilisation
15 data from data capture modules provided in the network
2 and at servers. Utilisation data will include
statistics such as server CPU (central processing
unit) utilisation, memory utilisation and network
bandwidth usage. Thus this information is of a more
20 generic data than that obtained directly from a user's
browser.

The database server 41 further comprises a
standard event and utilisation filter module 10 which
is used to process the data stored in the data log 9
25 in order to reduce its volume and present it in a more

usable fashion. The standard event and utilisation filter module 10 uses filter rules stored in a filter rule database 11.

5 A rule wizard system 12 is provided to allow a systems manager to generate filter rules. A filter rule is used to define standard events which are characterised by a sequence of one or more occurrences taking place in a given sequence within a single browsing session.

10 As described above, an occurrence is constituted by any action carried out by the user or any event within the browser however minor. A collection of data concerning such occurrences is not always useful. The generation of filter rules defining standard events
15 allows the data to be put in a much more usable and compact form. The rules wizard system 12 allows a newly generated filter rule to be tested against a portion of the data stored in the data log 9 to test its function. Once a rule has been tested it can be
20 stored in the filter rule database 11.

Each filter rule can have active or inactive status. This means that not all rules which have been generated must be applied to all incoming data. All data stored in the data log 9 is parsed by all active
25 rules specified in the standard event and utilisation

filter module 10.

Each filter rule assigns a name to each standard event and may also specify one or more associated value, such as an elapsed time, an absolute time, or
5 other data to be retrieved from the specific occurrences making up the standard event.

The standard event and utilisation filter module 10 holds a state table for each rule in an active browsing session. Whenever the rule state table
10 signifies that a standard event has occurred, a corresponding event record is written into a standard event log 13. The standard event record includes a time stamp taken from the final occurrence which triggered the standard event creation. Moreover, any
15 additional data specified by the rule will be inserted into the standard event record.

Rules may also be defined using the rule wizard system 12 to identify utilisation events. A state table is held in the standard event and utilisation
20 filter module 10 for each utilisation event rule. Whenever the utilisation event rule state table signifies that a utilisation event has occurred, a utilisation event record is written into a utilisation log 14.

25 The standard event and utilisation filter module

10 also has the function of calculating real time statistical information which it stores in a real time statistics database 15.

5 Data from the real time statistics database 15 is fed to a real time display module 16 which displays real time statistical data on the performance of websites which are being monitored. The type of information which is provided by the real time display module 16 is high level and thus, for example, might
10 include number of sales per second.

A graph generation module 17 is provided to generate graphs from the data stored in the standard event and utilisation log 13,14. These graphs may be made available via suitable web pages and can
15 represent historic or near real time information. A graph may for example show data concerning the different times taken to load a given page on a web site. The system includes a facility such that by appropriately selecting points on a graph, the
20 detailed standard event record may be viewed for that or those points. Furthermore, it is possible to display all occurrences which resulted in that standard event.

Tools provided in the systems maintenance module
25 7 as well as the functionality of the real time

display module 16 and graph generating module 17 allow the production of ad hoc reports by customers who are having their web sites monitored.

It is not essential for customers to gain
5 performance information directly from the central unit
4. Reports may be generated by the administrators of the monitoring system and provided to customers. Such reports may, for example, be in the form of hard copy or sent electronically.

10 The database server also includes a session replay module 18 which allows the replay of any session stored in the data log 9. Thus the functionality of the session replay module 18 is similar to, but distinct from, that supplied by the
15 Slayer system described above. The session replay module 18 functions by retrieving all the HTML code stored in the data log 9 and optionally loading any unstored data such as images from the Internet. The module 18 then replays the actions taken by the user.
20 This module gives the ability to review the actions carried out by a particular user which may be useful in identifying a problem which occurred either to assist the user in overcoming the problem or to highlight an aspect of the web page or web site which
25 may be improved.

- The systems described above are capable of providing information concerning uses of web sites which is extremely detailed. For example it is possible to provide real time displays and or reports
- 5 showing the following amongst other things:
- Number of visitors per unit time,
 - Number of new customers,
 - Number of customers who ordered,
 - Number of customers by location,
 - 10 Number of visitors who tried and failed to order,
 - Total website orders,
 - Total website order value,
 - Website navigation,
 - Usability of website,
 - 15 Apparent speed per customer,
 - Typical route through site,
 - Page hot spots,
 - Time to complete a standard event,
 - Events by time of day,
 - 20 Average and maximum/minimum time to complete transaction,
 - Page load time by page by time of day,
 - Time from loading to exit of page,
 - Exit route from page,
 - 25 Name of user by time of day,

IP address of user by time of day.

It is envisaged that such and other data will enable organisations running websites to optimise their services. In particular it is envisaged that it
5 may be possible to reduce cost of providing the service, maximise revenue/transaction rate, highlight and investigate failure modes, understand usability of site, detect unsatisfactory service, detect and correct site design issues, undertake detailed
10 analysis of user behaviour, evaluate success of investments in technology, target IT investment strategies for e-commerce applications.

The following gives an example of the kind of advantages which using a system of the current
15 application can give.

A website owner may be experiencing a large number of incomplete transactions and also retrieving complaints of poor performance from their customer research. By including the Prophet applet on each
20 relevant page of their website and configuring the system to retrieve time of day information on the entry and exit from each page as well as recording the exit routes from each page the following type of information may be gained. For example the time taken
25 by users to complete some key pages may be identified.

For example it may be determined that whilst 50% of users complete a page in less than 20 seconds nearly 20% take more than 2 minutes, whilst 20% never complete the transaction. By more careful analysis of the data it may become apparent that a large proportion of the slow users are using the same Internet service provider and this can indicate that there are problems using this Internet server provider. However, there may be some other reason whilst the other users are having difficulty. It might, for example, be established that a high proportion of users are having to access the help page more than twice whilst attempting to complete the transaction. This could point to a problem in the page design. Moreover, it might be established that a significant proportion of customers begin to fill out a user registration form but fail to complete it. This could suggest that the registration form is too long or too complex. It might alternatively be found that a large number of users leave the site via a banner advert provided on one of the pages. It may then be concluded that the positioning of the advert is disadvantageous.

As another example, it is highlighted that cursor/mouse position within a page can be monitored.

This can be used to generate statistical data on cursor position. This in turn can assist in page design, it might for example, be established that a majority of users seem to be attempting to find the
5 "submit" button in the bottom right hand corner of a page even though the button is located elsewhere.

In general terms it will be appreciated how the information provided by systems described in the present application can be of great assistance in
10 improving website design.

In alternatives a rules wizard may not be included but rather rules may be generated and programmed manually.

CLAIMS

1. A method for monitoring user interactions with a service provided over a network to a plurality of end user communications units comprising the steps of:
5 monitoring and recording activities instigated by at least one user when using the service by interacting with one or more electronic page displayed in a browser on a respective communications unit, the recorded activities including a log of pointing device
10 position on the or each page; and playing back activities recorded in the above step, wherein the playing back step includes redisplaying the electronic pages originally displayed and redisplaying user interaction therewith.
15
2. A method according to claim 1 comprising the step of receiving at a central unit, information obtained at each end user communications unit during monitoring of the respective user's interaction with the service.
20
3. A method according to claim 2 wherein the information is obtained at each end user communications unit under the control of respective applets running within the respective browsers;
25 and the method comprises, before the information

receiving step, the further steps of:

receiving, at the central unit, activation notifications from active applets; and

outputting, from the central unit, respective
5 configuration messages to each communications unit
from which activation notification has been received,
which messages are arranged to configure the
respective applets so as to control the information
sent to the central unit.

10

4. Apparatus for monitoring user interactions with a
service provided over a network to a plurality of end
user communications units comprising means for
monitoring and recording activities instigated by at
15 least one user when using the service by interacting
with one or more electronic page displayed in a
browser on a respective communications unit, the
recorded activities including a log of pointing device
position on the or each page; and
20 means for playing back the recorded activities by
redisplaying the electronic pages originally displayed
and redisplaying user interaction therewith.

5. A computer program comprising code portions which
25 when loaded and run on computer means cause the

computer means to execute the method of any one of
claims 1 to 3.

6. A computer program comprising code portions which
5 when loaded and run on computer means cause the
computer means to constitute the apparatus according
to claim 4.



INVESTOR IN PEOPLE

Application No: GB 0109290.7
Claims searched: 1-6

Examiner: Owen Wheeler
Date of search: 14 May 2001

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.S): H4P (PEUX)

Int Cl (Ed.7): G06F: 11/34; H04L: 12/24, 12/26, 29/06

Other: Online: EPODOC, JAPIO, WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	EP 1045547 A2 [NETZERO]	
A	EP 1026608 A2 [LUCENT]	
A	WO 00/55715 A1 [WEBCRITERIA]	
A	WO 99/13423 A1 [NARRATIVE]	
A	WO 98/43380 A2 [THINKING MEDIA]	

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.